Institute of Architecture of Application Systems

The vision for MUSE4Music. Applying the MUSE method in musicology.

Johanna Barzen[1], Uwe Breitenbücher[1], Linus Eusterbrock[2], Michael Falkenthal[1], Frank Hentschel[2], Frank Leymann[1],

1. Institute of Architecture of Application Systems, University of Stuttgart, Germany
{nachname}@iaas.uni-stuttgart.de

2. Musicological Institute, University of Cologne, Cologne, Germany
{Frank.Hentschel, Linus.Eusterbrock}@uni-koeln.de

BibTeX

```
@article {ART-2016-17,
    author = {Johanna Barzen and Uwe Breitenb{\"u}cher and Linus
Eusterbrock and Michael Falkenthal and Frank Hentschel and Frank Leymann},
    title = {{The vision for MUSE4Music. Applying the MUSE method in
musicology}},
    journal = {Computer Science - Research and Development},
    address = {Heidelberg},
    publisher = {Springer},
    pages = {1--6},
    type = {Artikel in Zeitschrift},
    month = {November},
    year = {2016},
    doi = {10.1007/s00450-016-0336-1},
    url = {http://www2.informatik.uni-stuttgart.de/cgi-
bin/NCSTRL/NCSTRL_view.pl?id=ART-2016-17&engl=}
 }
```


Universität Stuttgart
Germany

Noname manuscript No.
(will be inserted by the editor)

The Vision for MUSE4Music

Applying the MUSE Method in Musicology

Johanna Barzen · Uwe Breitenbücher · Linus Eusterbrock · Michael Falkenthal · Frank Hentschel · Frank Leymann

Received: date / Accepted: date

Abstract Investigating the emotional impact of historical music, e.g. music of the 19th century, is a complex challenge since the subjects that listened to this music and their emotions are forever gone. As a result, asking them for their experiences is not possible anymore and we need other means to gain insights into the expressive quality of music of this century. In this vision paper, we describe a pattern-based method called MUSE4Music to quantitatively find similarities in different pieces of music. The reconstruction of musical patterns will allow us to draw conclusions from erratic documents that go far beyond the single pieces they are referring to.

Johanna Barzen
Institute of Architecture of Application Systems
University of Stuttgart, Germany
E-mail: Johanna.Barzen@iaas.uni-stuttgart.de

Uwe Breitenbücher
Institute of Architecture of Application Systems
University of Stuttgart, Germany
E-mail: Uwe.Breitenbuecher@iaas.uni-stuttgart.de

Linus Eusterbrock
Musicological Institute
University of Cologne, Germany
E-mail: Linus.Eusterbrock@uni-koeln.de

Michael Falkenthal
Institute of Architecture of Application Systems
University of Stuttgart, Germany
E-mail: Michael.Falkenthal@iaas.uni-stuttgart.de

Frank Hentschel
Musicological Institute
University of Cologne, Germany
E-mail: Frank.Hentschel@uni-koeln.de

Frank Leymann
Institute of Architecture of Application Systems
University of Stuttgart, Germany
E-mail: Frank.Leymann@iaas.uni-stuttgart.de

Keywords Pattern Discovery · Musical Patterns · Automation · Mining · Digital Humanities · Musical Expressivity

1 Introduction, Background, and Motivation

If, as historians of music, we want to investigate the emotional impact of music — e.g. the expressive quality of music of the 19th century — we encounter a problem: the subjects that listened to this music and were seized by some emotion, moved to tears or beside themselves are long gone. We cannot ask those persons what they felt and what they liked about the music and why. While we can conduct experiments with living people about how they experience music (see, for instance, Zentner et al (2008)), we cannot do the same with, say, a 19th-century audience. Now, one might make the claim that the music of Beethoven, Chopin, Berlioz, Verdi, and the like is part of a living tradition essentially unaltered since the 19th century and that we may just as well investigate contemporary listeners. However, such a solution would never satisfy a historian who is all too aware of vast change that has taken place over the intervening years and difference between historical and contemporary cultural contexts. After the development of atonal music, Free Jazz, rhythm dominated rock music, and many other new listening experiences, it is hardly possible that we hear music of the 19th century with the same ears as people from that time did. And even if, in the end, we should learn that our experience is still essentially unchanged when compared to the experience of the audiences of the day, we cannot presuppose such an equivalence. Methodologically, this would be a dubious proposition.

What can we do to deal with this methodological challenge? And how might computer science help? By applying *computer-aided pattern research* to the interpretation of music, we are convinced that it is possible to gain valuable insights into the way listeners of the past responded to music. However, to explain the approach we envision, we need to take a step backwards and examine traditional research methods in the field. Historical research is necessarily dependent on documents from the time investigated. So, more precisely, the problem is that we lack sufficient primary sources that tell us how people in the 19th century were affected by music. Yet, there are a couple of documents and hints we may use, for example, reviews, autobiographies, letters, texts of vocal music, narratives in programme music, or musical quotations. However, they are rarely detailed enough, occur only sporadic, and usually refer to a specific composition so that it is impossible to make general statements about the effect on the recipient of a certain kind of music.

We are looking for a method that will allow us to draw conclusions from such scattered documents that go far beyond the single pieces they are referring to and the contexts they are rooted in — without simply disregarding the traditional toolkit of the historian. That search led the musicologists in our team to explore techniques similar to pattern research in other fields (specifically, in the sense of Christopher Alexanders architectural pattern (Alexander et al, 1977)). However, we only discovered this similarity after the fact. Our team introduced the concept of musical expression types ("*Ausdruckstypen*", see Hentschel (2013)). Our thought process was as follows: (i) if information about the emotional effect of a musical piece has been documented, and (ii) if this information can be related to a specific passage of that piece, and (iii) if finally this passage employs characteristic compositional techniques ("*musikalische Faktur*") that are not unique but have "relatives" in the 19th-century repertoire, it is possible — within a certain historical and cultural context (and with due caution) — to apply such knowledge about the emotional impact of one specific piece to the related pieces[1]. These pieces are related by the use of similar compositional techniques, i.e., by a *compositional pattern*. If such a pattern can be determined to correspond to a particular emotional meaning, we call it a *type of musical expression*. So, for example, if we find some pieces that contain passages that are identified as expressing joy by historical sources and if we can determine a pattern in the compositional techniques used in these pieces, we may assume that this pattern was used to express joy. We would call such a pattern expressing joy a type of musical expression and name it *The Joyous*.

However, finding compositional patterns and determining their expressive content presents musicology with serious problems: First, there is the sheer volume of musical pieces that has to be examined. The standard repertoire from the 19th century is but the tip of the iceberg. There are many different possible expressive types in one single piece, for instance *the martial*, *the dreadful*, *anger*, or *resoluteness*. Also, consider how easily we are biased in the interpretation of the emotional impact of historical music, for example because we associate the piece with a unique but maybe misleading document (like the famous fate that was said to knock at the door in Beethoven's 5th) or because we associate the music with our own personal memories that may be wholly unrelated to the music itself or simply because we naïvely confuse our emotional experiences with the historical. Such were the problems the musicologists participating in this vision paper (University of Cologne) were struggling with when they became aware of the *MUSE project*[2], which seeks to develop a method and accompanying tools for applying pattern research to media studies. This project is run by a team of computer scientists at the University of Stuttgart — the second group of authors of this vision paper. The *MUSE method* (as described in Barzen and Leymann (2014)) includes (i) the identification of relevant characteristics of the domain (e.g. film, theatre, or music) and structuring them by taxonomies and ontologies, (ii) a detailed capturing of the concrete solutions (i.e. concrete costumes in films or scores and audio files in music) by their relevant characteristics, (iii) analysing the "effect" on the recipient of each of the concrete solutions, and (iv) abstracting those with the same "effect" as an abstract solution. This abstract solution is described as a pattern (cf. Falkenthal et al (2014a) and Falkenthal et al (2014b)). The method is accompanied by a tool chain containing a solution repository and a pattern repository, which supports the process of pattern identification as described in Fehling et al (2014). The MUSE method and tool chain have already been applied to

[1] Within musicology there is a tradition of research on musical topics that began with a monograph by Ratner (1980). Other important contributions include the work of such authors as Kofi Agawu, Wye Allanbrook, Robert Hatten and Raymond Monelle (see McKay (2007) and Mirka (2014)). The basic idea of this area of research, which is almost exclusively confined to late 18th-century music, becomes clear in Danuta Mirka's definition of a musical topic as "*musical styles and genres taken out of their proper context and used in another one*" (Mirka (2014), p. 2). We draw on this approach. However, our concept of pattern will not be limited exclusively to styles that are taken out of their proper context but will extend to characteristic sets of musical properties in any context.

[2] http://www.iaas.uni-stuttgart.de/forschung/projects/MUSE/indexE.php

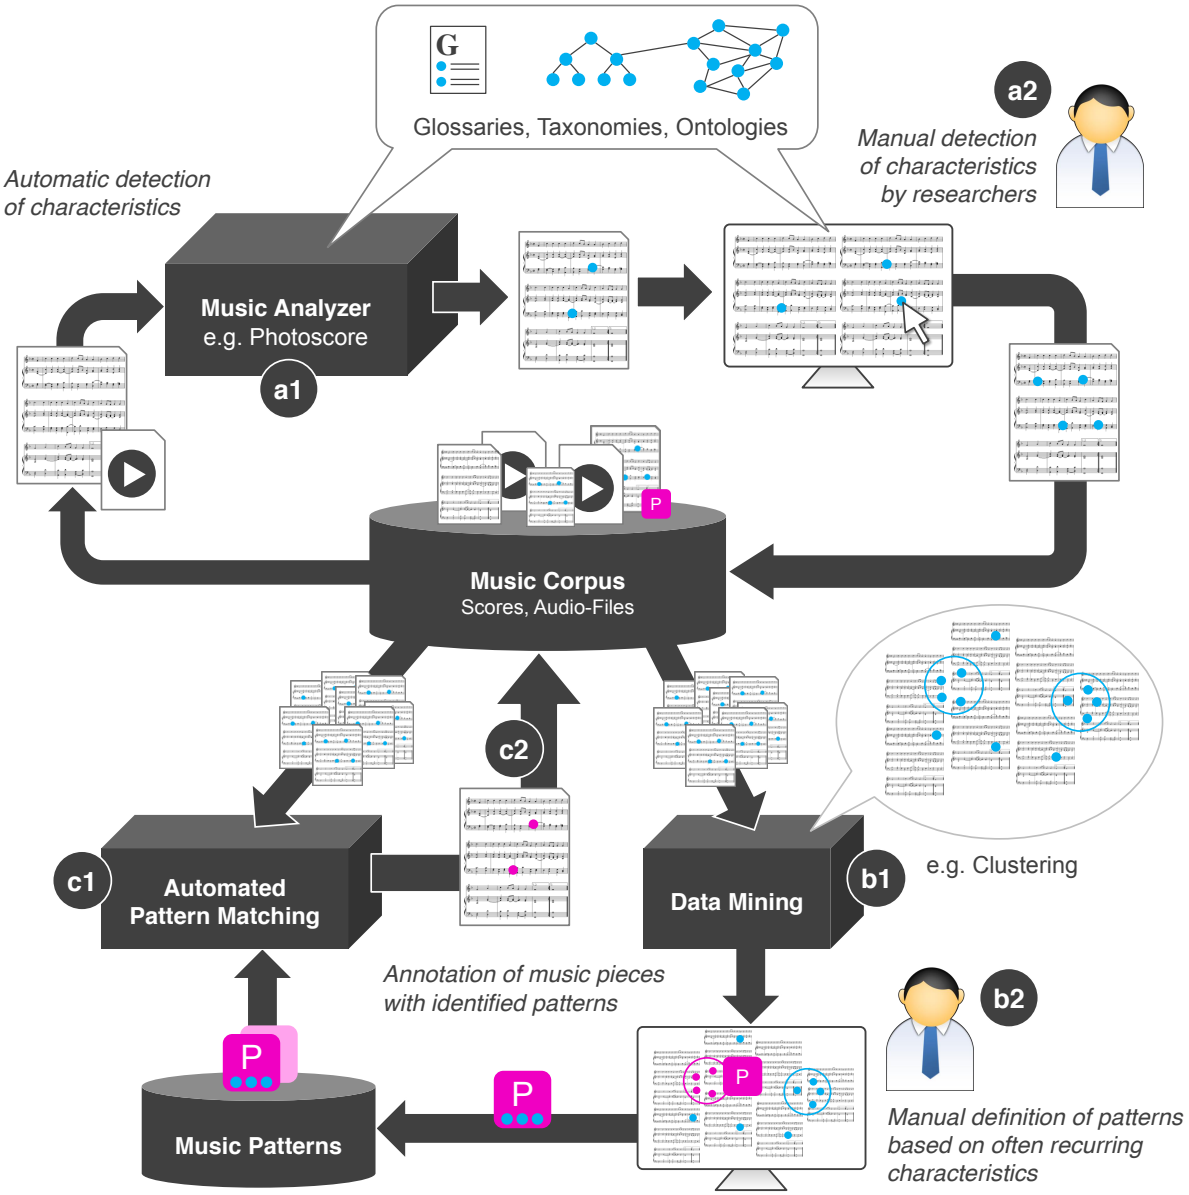

Fig. 1 The MUSE4Music Method

find *costume patterns* in film studies (Schumm et al, 2012). The approach seems generic enough to be applied to other fields of the humanities (cf. Barzen and Leymann (2016)), especially to the field of musicology. The corresponding application of the MUSE method to the problems sketched above is referred to as *MUSE4Music*, which is the vision we want to illustrate in this paper. Although the overall method of MUSE is applicable, the capturing of the concrete solutions as well as the analysis of these solutions need a lot of adaptation to the requirements of the special domain of musicology.

2 The MUSE4Music Vision

In this section, we present our vision for MUSE4Music — a method to support musicologists in historically interpreting facts to gain insights into the expressive quality of music of the 19th century. The major goals of MUSE4Music are (i) to discover *musical patterns* based on a semi-automated analysis of similarities in different pieces of music, (ii) to associate these patterns with a certain kind of emotion, e.g., joy or anger, and (iii) to automatically find the discovered patterns in other musical pieces to gain new insights about music that has never been analysed with respect to these emotions.

2.1 Overview of the MUSE4Music Method

An overview of the proposed method is shown in Figure 1. The method consists of three different phases, which are explained in detail in the following: (a) *detection of musical characteristics*, (b) *pattern discovery*, and (c) *pattern matching*. The phases are connected by a central *music corpus* that contains scores and audio files. In the first phase, this corpus is enriched by annotations, which provide the basis for discovering patterns in the second phase and, in the third phase, these patterns are identified in other pieces that have not been used before to discover these patterns.

In the first phase, scores and audio files are analysed in order to detect musical characteristics such as harmony, rhythm, melodic features, instrumentation, and dynamics. This analysis combines an automated step and a manual step: First, in step a1, a component called *music analyser* automatically consumes scores and (if available) corresponding audio files as input and analyses these pieces for characteristics that can be recognized, for example, rhythmic and tone sequences, or musical phrases. Because there are several tools available that promise to support the automated analysis of scores and audio files, e.g., Photoscore[3] or melodicmatch[4], the music analyser is built as a framework that supports the integration of different individual tools based on a plug-in mechanism. To provide a clear understanding of the detected characteristics, a glossary, taxonomies, and ontologies are used to structure the domain of music and to define the terminology and semantics of the annotated information. The outputs of this step a1 are scores that have been automatically annotated with musical characteristics. In step a2, musicologists manually annotate the resulting score with characteristics that cannot be detected automatically, such as complex harmonies. In addition, due to various discussions about the quality of automatically detected musical characteristics, in step a2, a musicologist may also check the automatically detected characteristics for accuracy and meaningfulness where necessary. Therefore, the first phase is semi-automated and provides the possibility for the musicologist to make qualitative improvements to the automatically detected characteristics. The output of the first phase is stored in the music corpus, which now contains scores, audio files, as well as annotated scores after this first phase.

In the second phase, the annotated files in the music corpus are analysed using a *data mining* component (b1). This component tries to find similarities in different annotated scores to detect characteristics that occur frequently, e.g., by using clustering techniques. Although this step produces quantitative results, the meaningfulness of the found similarities may be called into question. Therefore, a musicologist manually inspects the results in step b2 to decide which similarities have a certain effect and which of them are meaningless. Thus, the quantitatively analysed similarities are qualitatively verified. Those similarities that have a certain effect are defined as *music patterns*. But not all of these patterns lead to an answer to the question of what kind of expression those patterns may have and how they affected a contemporary audience. Therefore, due to the musicology research interests outlined in the introduction, the musicologists also collect and interpret historical documents that may lead to an understanding of what emotions have been aroused by those patterns. If such a pattern with its expressive qualities can be determined by this *historical reading*, it can be referred to as a type of musical expression. The music patterns, which are detected as a type of musical expression, are stored in a *music patterns* database. The musicologist therefore adds a textual description of the respective meaning to the pattern and defines the semantics between the patterns to create the pattern language. Moreover, the stored pattern also contains a formalized description of the important characteristics that make up the pattern in order to enable the automatic matching in the next phase.

In the third phase, an *automated pattern matching* component investigates new scores to identify instances of already detected patterns in these musical pieces (c1). This *pattern matching* is enabled by the formal description of the characteristics of each pattern. For example, if a pattern called *The Joyous* is detected and described by its relevant harmonic, rhythmic, melodic features, instrumentation, and dynamics these parameters can be matched to the annotation of other scores to identify other instances of *The Joyous* pattern. These identified parts of the scores are annotated with the matching patterns and are also stored in the music corpus (c2).

2.2 Pattern-based Music Analysis: The MUSE4Music Framework

The method presented in the previous section will be implemented as the *MUSE4Music Framework*, an IT-system that supports musicologists in (i) running the automated components, e.g., the music analyser, as well as (ii) in manually detecting characteristics, checking the patterns identified by the system, and interpreting the findings. Therefore, the framework will provide an

[3] http://www.sibelius.com/products/photoscore/ultimate.html (accessed on March 11th, 2016)

[4] http://www.melodicmatch.com/index.html (accessed on March 11th, 2016)

intuitive user interface that supports musicologists in annotating scores. In particular, the user interface will link scores to associated audio files to enable an efficient detection of those characteristics that are mainly recognized only when listening to music as opposed to just reading scores. Moreover, the framework will graphically depict patterns in scores that have been identified by the framework. The MUSE4Music Framework will, thus, support the historical reading of the expressive quality of music based on a huge amount of processed data, which is only possible due to the powerful IT-technologies employed such as data mining.

2.3 Benefits of MUSE4Music

Applying the MUSE method to music seems to be a perfect solution to the problems discussed in the introduction. The central music corpus not only allows us to handle a huge amount of music, but also prevents possible biases. In detail, the proposed MUSE4Music method supports the following benefits:

1. *Controlling the historians subjectivity*: Being interposed between sources and their interpretation, the computer will help to reduce the subjectivity of the research process (without, of course, being able to entirely eliminate it since hermeneutical steps are still needed as an integral part of the overall method). Moreover, a precise and systematic capturing of the detected characteristics of musical pieces is supported by a specially constructed user interface.
2. *Allowing complex queries*: Via the MUSE4Music Framework, we have access to a database in the form of the music corpus that allows complex queries. Thus, if the musicologist has a suspicion of the existence of a certain expressive type, he or she may check this assumption. Thus, the framework helps to determine the central features of the supposed type.
3. *Finding new candidates for expressive types via automated mining*: The MUSE4Music Framework will also search for salient clusters of characteristics that may hint at expressive types (of course, it is the task of the musicologist to decide whether those clusters are significant or not). With this function the framework not only helps to support or falsify assumptions but contributes to creating new valuable assumptions.
4. *Detecting determined patterns in a continuously growing corpus*: The MUSE4Music Framework will analyse newly added musical pieces automatically for the occurrence of patterns that have already been discovered based on the patterns characteristics. A manual analysis of these results can also support or falsifies the accuracy of patterns found in previous iterations of the process.

3 Conclusion

The MUSE4Music approach described in this vision paper combines methods from musicology and computer science. The heart of our method consists of a pattern discovery procedure and the interpretations of these patterns as musical expression types. Expression types play an important part in expressing emotions in the music contained in our corpus. Therefore, the MUSE4Music method promises to provide an invaluable methodological progress for the research on historical musical expressivity.

Nevertheless, our approach has certain limitations. One limitation lies in the huge amount of potential material. We have to begin with a manageable, but growing corpus. Another limitation is rooted in the still rather limited capability of the automated music analysis tools. However, combining these tools seems to be a promising way to overcome these difficulties and support the annotations of the scores.

References

Alexander C, Ishikawa S, Silverstein M (1977) A Pattern Language: Towns, Buildings, Construction. Oxford University Press, New York

Barzen J, Leymann F (2014) Costume languages as pattern languages. In: Proceedings of Pursuit of Pattern Languages for Societal Change (PURPLSOC) - Preparatory Workshop, epubli GmbH, pp 88–117

Barzen J, Leymann F (2016) Patterns as Formulas: Applying the Scientific Method to the Humanities. Tech. Rep. 2016/01, University of Stuttgart

Falkenthal M, Barzen J, Breitenbücher U, Fehling C, Leymann F (2014a) Efficient Pattern Application: Validating the Concept of Solution Implementations in Different Domains. International Journal On Advances in Software 7(3&4):710–726

Falkenthal M, Barzen J, Breitenbücher U, Fehling C, Leymann F (2014b) From Pattern Languages to Solution Implementations. In: Proceedings of the 6$^{\text{th}}$ International Conferences on Pervasive Patterns and Applications (PATTERNS), pp 12–21

Fehling C, Barzen J, Falkenthal M, Leymann F (2014) PatternPedia - Collaborative Pattern Identification and Authoring. In: Proceedings of Pursuit of Pattern Languages for Societal Change (PURPLSOC) - Preparatory Workshop, epubli GmbH, pp 252–284

Hentschel F (2013) Festlichkeit: Expressive Qualität und historische Semantik bei Beethoven. Archiv für Musikwissenschaft 70 pp 161–190

McKay N (2007) On topics today. Zeitschrift der Gesellschaft für Musiktheorie 4(1):159–183

Mirka D (2014) The Oxford Handbook of Topic Theory. Oxford University Press, New York

Ratner L (1980) Classic Music: Expression, Form, and Style. Schirmer, New York

Schumm D, Barzen J, Leymann F, Ellrich L (2012) A Pattern Language for Costumes in Films. In: Proceedings of the 17$^{\text{th}}$ European Conference on Pattern Languages of Programs (EuroPLoP), ACM, pp 221–249

Zentner M, Grandjean D, Scherer KR (2008) Emotions evoked by the sound of music: Characterization, classification, and measurement. Emotion 8(4):494–521